

Remarks

Entry of the amendments, reconsideration of the application, as amended, and allowance of all pending claims are respectfully requested. Upon entry of the amendments, claims 50-97 are pending. **Applicants respectfully request that these Remarks be carefully considered by the Examiner.**

In an effort to advance prosecution of this application, and not in acquiescence to any rejections, applicants have amended the claims to more particularly point out and clearly define their invention. Specifically, applicants have cancelled claims 1-49 and have added claims 50-97. Support for the amendments can be found throughout the specification, and therefore, no new matter is added. For example, support for claims 50-58, 66-74 and 82-90 can be found on pages 14, and 16-20 of applicants' specification and in FIGs. 9a-9b; support for claims 59-62, 75-78 and 91-94 can be found on pages 14-16 of applicants' specification and in FIGs. 6b-6d; and support for claims 63-65, 79-81 and 95-97 can be found on pages 20-21 of applicants' specification and in FIG. 15. Again, applicants have rewritten the claims, in an effort to advance prosecution of this application.

In the Office Action, dated April 26, 2005, claims 1-49 are rejected under 35 U.S.C. 102(e) as being anticipated by Shrivastava et al. (U.S. Patent No. 6,449,734 B1). Applicants respectfully, but most strenuously, traverse this rejection for the reasons below to any extent deemed applicable to the new claims.

Applicants' invention is directed to managing processing groups of a distributed computing environment. In one aspect, applicants' invention is directed to a protocol used to join a prospective member to a processing group. As one example, the protocol is used to join a prospective member to an active processing group. The join protocol includes various steps taken to ensure configuration consistency. Specifically, a sequence number is used throughout the join protocol to control whether a prospective member joins the group and to ensure configuration consistency.

As one particular example, applicants claim a method of managing processing groups of a distributed computing environment (e.g., independent claim 50). The method includes, for

instance, requesting via a request by a prospective member to join a processing group of the distributed computing environment, the request including a sequence number indicating a version of the processing group; determining whether the prospective member can join the processing group, the determining employing the sequence number; and joining the processing group by the prospective member, in response to at least in part to the determining indicating that the prospective member can join the processing group. Thus, in this aspect of applicants' claimed invention, a prospective member requests to join a processing group and that request includes a sequence number indicating a version of the processing group. That sequence number is also used in determining whether the prospective member can join the processing group. This is very different from the teachings of Shrivastava.

While applicants agree that both Shrivastava and applicants' claimed invention use sequence numbers, applicants' use of sequence numbers is very different from that of Shrivastava. Thus, applicants' claimed invention is not described, taught or suggested by Shrivastava.

For example, applicants claim requesting via a request by a prospective member to join a processing group of the distributed computing environment, the request including a sequence number indicating a version of the processing group. In this aspect of applicants' claimed invention, a sequence number is included in a join request, and that sequence number indicates a version of the processing group. There is no description in Shrivastava of a join request having a sequence number indicating a version of the processing group, as claimed by applicants. Instead, in Shrivastava, to join an active cluster, an authentication sequence is performed, and if it is successful, the database of the arriving node is updated, if necessary (Col. 11, lines 40-50). There is no discussion in the join protocol of Shrivastava of requesting to join via a request that includes a sequence number indicating a version of the processing group, as explicitly claimed by applicants. For at least this reason, applicants respectfully submit that their claimed invention is not anticipated by Shrivastava.

As a further example, applicants claim determining whether the prospective member can join the processing group, in which the determining employs the sequence number provided in the request. Again, this is not described, taught or suggested in Shrivastava. In Shrivastava, a

node joins the group after an authentication sequence is performed and then the database is updated. There is no discussion in Shrivastava in the joining process of employing a sequence number to determine whether the node can join the group. Since this is missing from Shrivastava, applicants respectfully submit for at least this reason, applicants' claimed invention is not anticipated by Shrivastava.

Again, although Shrivastava uses sequence numbers, the use of the sequence numbers in Shrivastava is different from that claimed by applicants.

In Shrivastava, sequence numbers are attached to transactions which are logged. If two transactions in the log have the same sequence number, then the first transaction is discarded. Thus, the sequence numbers are used to maintain consistency between a new cluster and a previous cluster by ensuring that transactions that failed to be replicated on all the members of previous clusters are discarded when forming the new cluster. The sequence numbers of Shrivastava are not used in a request to join and are not used to determine whether a prospective member can join a cluster.

To further explain, in one aspect, applicants are claiming a particular technique for a prospective member to join a processing group. This specific technique uses sequence numbers as an integral part of the technique. In contrast, while Shrivastava also includes a technique for joining a cluster, its technique does not make the same type of decisions based on a sequence number, as claimed by applicants. In Shrivastava, the join technique includes configuring and mounting devices; communicating to the last known members of the cluster; if a member is discovered, performing an authentication sequence; and if authentication is successful, examining the database in the arriving node, and if it is out of date, then providing an updated copy of the database (Col. 11, lines 30-50). There is no discussion in Shrivastava's technique of requesting via a request to join a processing group in which that request includes a sequence number indicating a version of the processing group. Further, there is no description in that technique of determining whether the prospective member can join the processing group in which that determining employs the sequence number. Thus, applicants respectfully submit that Shrivastava teaches a different join technique than that claimed by applicants.

In addition to the above, Shrivastava also describes forming a new cluster. In Shrivastava, in order to form a new cluster, the member configures and mounts devices; communicates to the last known members of the cluster; if there is no discovery of the members, it gains access to a quorum device that has a log file of transactions and unrolls the log file (see, e.g., Co. 11, lines 50-65). It is in the unrolling of the log file, that the sequence numbers are used to determine what data is to be unrolled. If transactions logged in the log file have duplicate sequence numbers, then the first transaction with that number is discarded. However, again there is no description, teaching or suggestion of requesting via a request by a prospective member to join a processing group in which the request includes a sequence number indicating a version of the processing group. Again, there is no description, teaching or suggestion of determining whether the prospective member can join the processing group and that determining employing a sequence number, as claimed by applicants.

For at least the above reasons, applicants respectfully submit that independent claim 50, as well as independent claims 66 and 82, is patentable over Shrivastava. Further, independent claims 58, 74 and 90 are also patentable over Shrivastava for the same reasons as above, as well as for their own additional features.

For example, independent claim 58 explicitly recites in the join protocol that the prospective member compares the sequence number in the join request with a current group sequence number. Since Shrivastava does not describe, teach or suggest a sequence number provided in a join request, it follows that Shrivastava fails to describe, teach or suggest comparing the sequence number in the join request with the current group sequence number, as claimed by applicants. Instead, the comparisons of sequence numbers in Shrivastava occur when a new cluster is being formed and a transaction log is unrolled. In that case, if two transaction numbers have the same sequence number, then one of the transactions is discarded. It does not affect whether the new cluster is formed.

As a further example, the join protocol explicitly claimed in independent claim 58 specifically recites updating the current group sequence number as part of the join protocol. Again, this is not described, taught or suggested in Shrivastava. Instead, in Shrivastava the

sequence number is updated when a new transaction is logged. There is no description, teaching or suggestion of increasing the sequence number in the join protocol of Shrivastava.

For at least the above reasons, applicants respectfully submit that independent claims 58, 74 and 90 are patentable over Shrivastava.

In a further aspect of applicants' invention, applicants claim a method of managing processing groups of a distributed computing environment (e.g., independent claim 63). The method includes, for example, detecting a failure of a member of a processing group of the distributed computing environment; and excluding the failed member from the processing group in response to the detecting, wherein the excluding includes updating a sequence number of the processing group to exclude the failed member, the sequence number identifying a version of the processing group. Thus, in this aspect of applicants' claimed invention, in order to exclude a failed member, a sequence number is updated. This is not described, taught or suggested in Shrivastava.

Shrivastava does describe that a failing node is removed from a cluster; however, it does not describe applicants' claimed technique for removing the member from the cluster. In Shrivastava, if a node fails to acknowledge an update, then that failing node is removed. Further, if a node fails to commit a transaction, then again, that node is removed from the cluster. This is described in Col. 7, lines 30-32 and lines 55-60 of Shrivastava. In particular, in Shrivastava, when a failure is detected by a system of the cluster, the detecting system broadcasts a message to the cluster causing other members to verify their view of the current cluster membership. If a system does not respond, then it is removed from the cluster (e.g., Col. 5, lines 45-55). There is no description in the technique of Shrivastava of detecting a failure of a member of the processing group, and then excluding that member via updating a sequence of the processing group, as claimed by applicants. Since this is missing from Shrivastava, applicants respectfully submit that Shrivastava does not describe, teach or suggest applicants' claimed invention.

For all of the above reasons, applicants respectfully submit that their invention, as explicitly claimed in the claims, is patentable over Shrivastava. The dependent claims are patentable for the same reasons as the independent claims, as well as for their own additional

features. Thus, applicants respectfully request an indication of allowability for all pending claims.

Should the Examiner wish to discuss this case with applicants' attorney, please contact applicants' attorney at the below listed number.

Respectfully submitted,

Blanche E. Schiller
Blanche E. Schiller
Attorney for Applicants
Registration No.: 35,670

Dated: July 26, 2005.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579